

Getting More From Your Multicore: Exploiting OpenMP From An Open Source Numerical Scripting Language

Michael S. Noble

Kavli Institute for Astrophysics, Massachusetts Institute of Technology

June 26, 2007

We introduce SLIRP, a module generator for the S-Lang numerical scripting language, with a focus on its vectorization capabilities. We demonstrate how both SLIRP and S-Lang were easily adapted to exploit the inherent parallelism of high-level mathematical languages with OpenMP, allowing general users to employ tightly-coupled multiprocessors in scriptable research calculations while requiring no special knowledge of parallel programming. Motivated by examples in the ISIS astrophysical modeling & analysis tool, performance figures are presented for several machine and compiler configurations, demonstrating beneficial speedups for real-world operations.

Categories and Subject Descriptors: D.3.2 [**Programming Languages**]: Language Classifications—*Very high-level languages, Concurrent, distributed, and parallel languages*; D.3.4 [**Programming Languages**]: Processors—*Code generation*; D.2.8 [**Software Engineering**]: Metrics—*Performance measures*

Additional Key Words and Phrases: Scientific Computation, Interactive Analysis, Astrophysics

1. INTRODUCTION

We are witnessing the arrival of serious multiprocessing capability on the desktop. While single-user machines with multiple CPUs have been available for several years, they remain uncommon and have typically only doubled, or to a lesser extent quadrupled, the number of processors. That is changing, though, as multicore chip designs begin to make it economical for typical users to access 8, 16, or even more processor cores. There is widespread concern, however, that it will not be easy to harness all of this power *within* applications on the desktop [Creager, 2005; Chuang et al., 2006]. Most desktop software has been developed for one CPU, and writing multithreaded software has traditionally been difficult [Lee, 2006].

Drawing from our own scientific niche, we recently noted how rarely parallel computing is employed for common modeling and analysis computations in observational astrophysics [Noble et al., 2006]. Researchers in other fields indicate a similarly low adoption of parallel methods by general investigators in their disciplines, e.g. [Creel, 2005]. Moreover, even if parallel programming were “easier,” the large bodies of serial software developed in communities over decades, and the mindsets they embody, cannot be changed overnight. Another difficulty is maintaining trust in such codes – instilled by years of vetting through the process of scientific publication – as they are retrofitted for parallelism. It has therefore been easier and safer for general practitioners to increase performance by purchasing faster serial hardware, rather than revamping algorithms or techniques for parallelism. Chip manufacturers are effectively telling us with multicore designs that

this tactic will not remain viable for much longer.

1.1 Very High Level Numerical Languages

At the same time researchers are well versed in scripting, particularly with array-oriented numerical languages like MatLab, Octave, and S-Lang, to name just a few. A key feature of these languages is that they allow easy manipulation of mathematical structures of arbitrary dimension, combining the brevity and convenience of an interpreted environment with most of the performance of compiled code. Operators and functions defined in the language of implementation (e.g. C) to work only with scalars are extended to interpreted arrays in the natural way, facilitating concise expressions such as `c = sin(a^3) + b*10.0` without regard to whether `a` or `b` are scalars, vectors, or multidimensional arrays. The high performance stems from moving array traversals out of the interpreted layer and into lower-level code like this fragment of C which provides vectorized multiplication in S-Lang:

```
case SLANG_TIMES:
    ...
    for (n = 0; n < na; n++)
        c[n] = a[n] * b[n];
    ...
```

One of the earliest motivations for SLIRP, the module generator for S-Lang, was to extend the reach of this vectorization to external C/C++ and Fortran codes.

1.2 OpenMP

The code above suggests that much of the strength and appeal of numerical scripting languages stems from relatively simple internal loops over regular structures. Another advantage of having these regular loops in lower-level compiled codes is that they are ripe for parallelization with OpenMP, a hardware-neutral specification aimed at facilitating parallel programming on shared memory multiprocessors. Conformant implementations of OpenMP offer a set of compiler directives for C/C++ or Fortran programs, supporting libraries, and environment variables which tune their operation. Programs are parallelized with OpenMP by tagging regions of code with comments in Fortran or preprocessor directives in C/C++. Proponents contend that conceptual simplicity makes OpenMP more approachable than other parallel programming models, e.g. message-passing in MPI or PVM, and emphasize the added benefit of allowing single bodies of code to be used for both serial and parallel execution. For instance, changing the above loop to

```
#pragma omp parallel for
for (n = 0; n < na; n++)
    c[n] = a[n] * b[n];
```

parallelizes the S-Lang multiplication operator; the pragma is simply ignored by a non-conformant compiler, resulting in a sequential program. OpenMP runtime environments have also been used for distributed computation over networked clusters, but this is not yet within the standard. Despite the promise of straightforward parallelism, the spread of OpenMP beyond high performance computing research groups has in part been hindered by the need for special compiler support. With a few exceptions such as OdinMP [Karlsson and Brorsson, 2004] and Omni [Kusano et al., 2000], this has come largely in the form of commercial compilers. The

shortfall of free compiler support for OpenMP helps explain the relatively small presence of OpenMP in open source numerical software, and marks the availability of OpenMP in GCC as a significant step towards the wider adoption of parallel computing by general practitioners.

1.3 Testbed

Our work was conducted primarily on 2 machine configurations: a dual-CPU (1.8 Ghz) Athlon workstation with 2 GB RAM running Debian 3.1 GNU/Linux, and 4 of 8 CPUs (750 Mhz) on a Solaris 5.9 server with 32 GB Ram. We refer to these as Linux2 and Solaris4, and use Linux1 and Solaris1 to denote serial execution. The author was the only user of Linux2 for all tests, while Solaris4 was shared with users running various jobs, many compute-intensive. Versions 1.9.3 and 2.0.7 of SLIRP and S-Lang were used, with codes executed in `slsh` and version 1.4.7 of ISIS [Houck, 2002], an astrophysical modeling and analysis tool developed at MIT.

2. RELATED WORK

We now highlight some of the more prominent efforts which relate to the three main areas encompassed by this paper: high-level numerical scripting environments, simplifying the use of tightly-coupled multiprocessors within them, and wrapper generators for them. In the commercial arena MatLab and Mathematica are among the most widely used high-level numerical environments to provide simplified support for multiprocessing. The MatLab*P extension strives to make it easy to adapt MatLab scripts for parallelism, by transparently storing arrays on, and offloading computations upon them to, a distributed compute server. IDL is arguably the most popular high level language in our astrophysics community; it also makes the use of multiple CPUs transparent, through an internal thread pool. The clear contrasts between these efforts and our work are cost and openness: commercial packages, while at times preceding their open source equivalents or eclipsing them in features, can be expensive to purchase and maintain. In addition, many research projects require software customizations which might either violate the proprietary nature of closed source or introduce unacceptable delays while waiting for fulfillment by vendors. Such conditions generally lead to the embrace of open methods.

In the realm of open software, OctaveHPC¹ was created to generalize Octave, a MatLab clone, for 64-bit platforms and integrate OpenMP directives so as to make transparent use of multiprocessors. The 64-bit extensions have been folded in to the source repository, but as of version 2.9.10 no OpenMP constructs appear within the Octave codebase nor do any papers appear to have been published describing the work. OdinMP [Karlsson and Brorsson, 2004] and Omni [Kusano et al., 2000] are notable for being among the earliest open-source compilers supporting OpenMP. Oriented towards academic research, neither has penetrated the wider open source community to nearly the same extent as GCC, nor do they appear to be actively maintained. The cOMPunity group² maintains a list of additional free OpenMP tools. SWIG [Beazley, 1996] is arguably the most powerful and widely used wrapper generator in the world. It can generate bindings to numerous scripting languages

¹<http://www.hpc.unsw.edu.au/OctaveHPC.html>

²<http://www.compunity.org>

and has considerably deeper support for C++ than does SLIRP. Unlike SLIRP, SWIG will not wrap Fortran, nor does it generate vectorized wrappers for numerical languages. The MatWrap³ tool has been used to generate vectorized bindings to the MatLab, Tela, and Octave matrix languages. Vectorization in MatWrap is not as advanced as in SLIRP (e.g. no support for Fortran, strings, or parallelization; arrays of unlike dimension may not be mixed), and it has not been maintained since 2001. The Tcl, Perl, and Python scripting languages are not vectorized, but actively-maintained numerical extensions do exist for each: BLT & NAP for TCL, PDL for Perl, and Numeric, NumArray, & NumPy for Python. Although some work has been done to create a SWIG extension which takes advantage of Python numerical extensions, it is not clear from the literature that vector-parallel wrappers can be auto-generated for these languages, particularly with the ease of SLIRP.

3. S-LANG

S-Lang is an extensible, C-like scripting language used in a number of popular free software packages and bundled with every major Linux distribution. It is highly suitable for scientific and engineering computation, offering a wide selection of extension modules and multidimensional numerics on par with commercial packages like MatLab and IDL. While a comprehensive analysis of the numerical engine of S-Lang is beyond the scope of this paper, Figs. 1 - 3 show 2 representative calculations on Linux1 which give a rough indication of its performance – in terms of speed and memory utilization – relative to IDL and the Tcl, Perl, and Python numerical extensions. With the exception of IDL, which we used from a (presumably optimized) binary installation, all codes were compiled with GCC 3.3.5 using -O3 and -fPIC flags. A total of 31 datapoints were gathered per language per computation, each representing the mean time of 1000 invocations of the respective calculation with a given array size (from 1 to 1e6 elements), using 8-byte real values. Memory statistics were collected by the `proc-time` utility,⁴ with small modifications.

4. SLIRP

SLIRP is a wrapper generator which simplifies the creation of dynamically loadable modules for S-Lang. It can dramatically reduce the effort needed to make external C/C++ and Fortran codes callable directly from the S-Lang interpreter. SLIRP is implemented in S-Lang and a small amount of supporting C code, with no other external dependencies, making it easy to install, use, and rapidly evolve. It has been used to generate bindings to a wide range of software, from libraries as large as Gtk and HDF5 to codes as small as numerical models for ISIS contained within a single file. It is particularly useful as a means of quickly assessing whether a piece of software is a good candidate module for ISIS; our group has evaluated numerous codes in this manner.

4.1 Vectorization

Perhaps the most distinguishing feature of SLIRP is its ability to vectorize wrapped functions, optionally tuned for parallelization with OpenMP. A stark example of

³<http://freshmeat.net/projects/matwrap>

⁴<http://www.xs4all.nl/%7Ejantien/software>

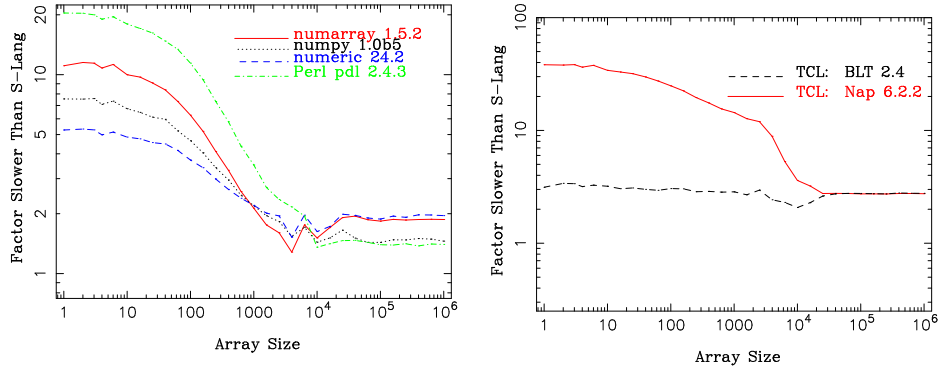


Fig. 1. Linux1 performance of Perl, Python, and Tcl numerical extensions, relative to S-Lang, on $\sqrt{b^2 - 4ac}$, where a, b, and c are arrays; smaller numbers are better.

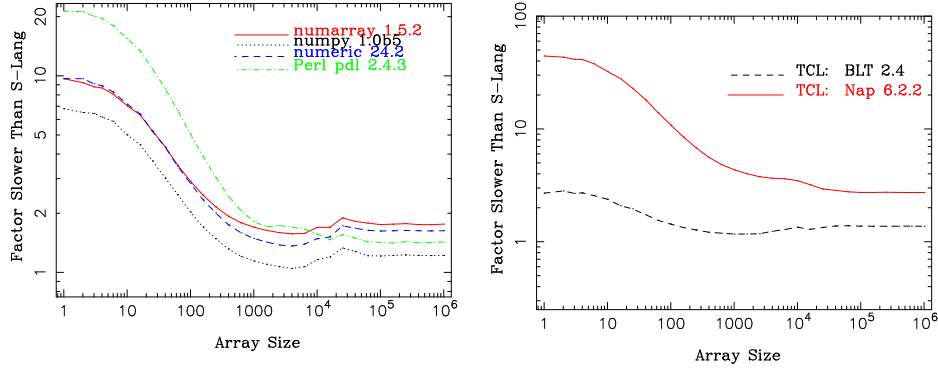


Fig. 2. Linux1 relative performance on the array slicing expression $(a^{1.5}/2 + b^2/4 + n^2 * \sin(c^3)/5) \leq n^2$, where n is the array size; smaller numbers are better.

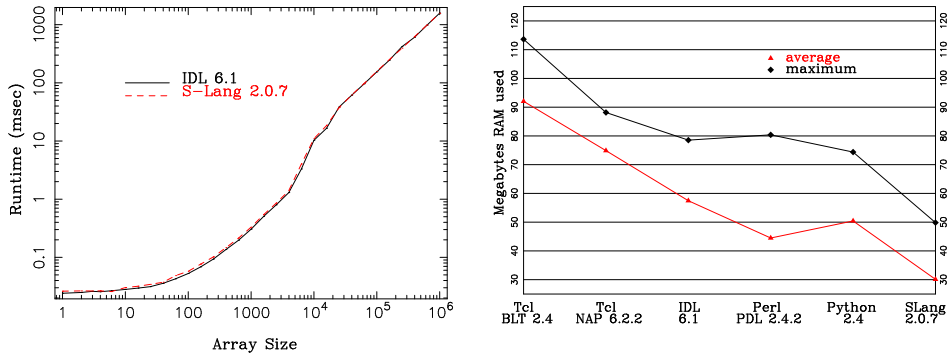


Fig. 3. Left: Linux1 performance of IDL 6.1 (binary) and S-Lang (statically linked), for $\sqrt{b^2 - 4ac}$. Right: memory consumed within each language environment when computing $\sqrt{b^2 - 4ac}$; smaller numbers are better.

```

isis> avol = array_map(String_Type, &sprintf, "%d", [1:100*100*100])
isis> tic; dvol = array_map(Double_Type, &atof, avol); toc
13.754

isis> import("atof")
isis> tic; pdvol = atof(avol); toc
0.1442

```

Fig. 4. Linux1 snapshot of S-Lang `atof` versus SLIRP vector-parallel version, on 100^3 strings.

the power of vectorization is given in Fig. 4, which was motivated by the desire to visualize a $320 \times 320 \times 320$ cube⁵ representing Doppler velocity mappings of Silicon II infrared emission observed with the Spitzer telescope. The 130Mb volume was supplied in ASCII form, so to minimize I/O time during exploratory analysis we first converted it to the high-performance HDF5 binary format, which involved some 320^3 calls to `atof`. This function is not vectorized in S-Lang, so to boost performance we generated a vector-parallel replacement, in only seconds with `slirp -make -openmp atof.h && make`, using SLIRP to also generate the Makefile. With faked data scaled down for didactic expedience to contain only 100^3 voxels, Fig. 4 indicates that the vector-parallel version is ca. 95X faster on our Linux2 machine. It is worth noting that `array_map()` is normally the fastest technique by which non-vectorized S-Lang intrinsics can be applied to arrays. Using other iterative mechanisms, say a `for` loop, would have yielded an even greater advantage for the parallelized `atof`. This snapshot also hints at another significant benefit of vectorization, namely brevity of end-user code. As a stronger example, consider the problem of reading N datasets from an HDF5 file [Folk et al., 1999], where each dataset is a 100-element single precision floating point array whose name is stored at index *i* of a string array. The minimal user-level code to read these data into a 2D array in IDL 6.1, without resource leaks, is given in the left panel of Fig. 5; the same result can be achieved with our vectorized HDF5 module in a single statement.

```

array = fltarr(N, 100)
fp = H5F_OPEN(file)
  for i = 0, N-1 do begin
    dp = H5D_OPEN(f, datasets[i])
    array[i, *] = H5D_READ(dp)
    H5D_CLOSE(dp)
  endfor
H5F_CLOSE(fp)

```

```
array = h5_read(file, datasets);
```

Fig. 5. Reading multiple HDF5 datasets with IDL 6.1 (left) and the vectorized S-Lang module.

Vectorization encompasses more than the simple promotion of scalar arguments to arrays. More generally, we say a function is vectorized when its arguments may be of multiple ranks. No distinction is made between the promotion of a rank

⁵See `volview` at <http://space.mit.edu/hydra/implement.html>

<pre>isis> vmult([1,2,3], [3,4]) Array shape or length mismatch isis> print(vmult([1,2,3], [5,5,5])) 5 10 15</pre>	<pre>isis> Arr = Double_Type[2,3] isis> Arr[0,*] = 5 isis> Arr[1,*] = 100 isis> print(vmult(Arr, [3, 4, 5])) 15 20 25 300 400 500</pre>
--	---

Fig. 6. Invoking a vectorized function with arrays of both similar and dissimilar shapes.

0 scalar to 1D or higher, a 2D array to 3D, and so forth. When a vectorized function is invoked with any argument whose rank exceeds that of its prescribed usage we say that both the argument and the function call are *vectorized*. SLIRP aims for maximum flexibility, allowing vectorized functions to be invoked using either scalar or array semantics and with few restrictions on the quantity, datatype, or dimensionality of arguments. For example, Fig. 6 shows a call mixing non-isomorphic arrays. The wrapped C function is prototyped as

```
void vmult(double *x, double *y, double *result, int len);
```

to multiply 2 vectors of length `len` and is called from S-Lang as

```
double[] = vmult(double[], double[])
```

The two signatures differ because an *annotation*⁶ has been applied to make the S-Lang usage more natural: **result** is moved from the parameter list to become a return value, and the vector length parameter is omitted because it can be obtained by inspecting the S-Lang arrays. The first `vmult` call fails for the obvious reason that vectors of dissimilar length cannot be multiplied. The second call succeeds, but is not vectorized because the ranks of both arguments match those of their prototyped parameters. The final call is vectorized because the rank of the first argument, a 2D array, exceeds its prototyped dimensionality of 1.

4.2 Dimensionality Theory

SLIRP uses a few simple metrics to decide whether a wrapper has been called with vectorized semantics, collectively referred to as the *parameters of vectorization*. To begin, each argument passed to a wrapper has an *expected rank*: a non-negative integer indicating the number of indices required to uniquely identify a single element. This rank is inferred at code generation time from the arguments signature within the function declaration. SLIRP distinguishes dimensioned arrays such as `double x[3][5]` from arrays of pointers like `double **x`, assigning them ranks of 2 and 1, respectively.

The *actual rank* of an argument is its dimensionality as passed at runtime. When the actual rank of any argument exceeds its expected rank, SLIRP needs to determine how many times the wrapped function should be called, or the *number of iterations* of the vectorization. This is decided by selecting a master array *M* – the input argument of highest rank – and computing the product of its excess dimensions. For example, if `Arr` in Fig. 6 was 4x3x3 instead of 2x3 then `vmult` proper would be called 12 times instead of 2. Formally, if *A* and *E* represent the actual

⁶Annotations are similar to SWIG *typemaps*, and are described in the SLIRP documentation.

and expected ranks of \mathbf{M} , and \mathbf{D} is a vector of length \mathbf{A} describing the size of each dimension of \mathbf{M} (in row-major form), then

$$Num_Iterations = \begin{cases} 1 & \text{when } \mathbf{A} = \mathbf{E} \\ \prod_{i=1}^{A-E} D[i] & \text{when } \mathbf{A} > \mathbf{E}. \end{cases} \quad (1)$$

Finally, SLIRP determines what to pass to the wrapped function by calculating a *stride* for each argument; this indicates by how much an index into the argument – viewed as a linear sequence of contiguous elements – should be advanced after iteration of the vectorization loop. Returning to the `vmult` call in Fig. 6, the strides of the first and second arguments are 3 and 0; within the wrapper the input arguments and return value are effectively represented as

```
double *arg1 = {5, 5, 5, 100, 100, 100};
double *arg2 = {3, 4, 5};
double *retval = malloc( sizeof(double) * 6);
```

and the 2 calls to `vmult` proper are executed as

```
vmult(arg1, arg2, retval, 3);
vmult(arg1+3, arg2+0, retval+3, 3);
```

Formally, the stride of \mathbf{M} and all isomorphic arguments is the number of elements contained within its expected dimensions

$$Stride = \prod_{i=A-E+1}^A D[i]. \quad (2)$$

The stride can be computed directly from the number of iterations by recalling that the number of elements in \mathbf{M} is the product of its dimensions

$$Num_Elements = \prod_{i=1}^A D[i]. \quad (3)$$

Factoring the left side into the product of excess and expected dimensions gives

$$Num_Elements = \prod_{i=1}^{A-E} D[i] \prod_{i=A-E+1}^A D[i], \quad (4)$$

and by noting that the first term here is the number of iterations we see

$$Stride = Num_Elements / Num_Iterations. \quad (5)$$

Equations 1 and 5 are coded into the `vec_pop()` routine discussed in the next section. Arguments not isomorphic to \mathbf{M} are legal as long as their number of elements equals the stride of \mathbf{M} ; they will be assigned a stride of 0.

4.3 Anatomy of a Vectorized Wrapper

To give a sense of what vector-parallelism entails, Fig. 7 shows the code generated for vectorized and parallel wrappers of the C `hypot` function. In the vectorized wrapper the return value and arguments of `hypot` are pointers, instead of scalars as

(a)

```

static void sl_hypot (void)
{
    double* retval;
    double* arg1;
    Slirp_Ref *arg1_r = ref_new(SLANG_DOUBLE_TYPE, sizeof(double), &arg1, 0x0);
    double* arg2;
    Slirp_Ref *arg2_r = ref_new(SLANG_DOUBLE_TYPE, sizeof(double), &arg2, 0x0);
    unsigned int _viter;
    VecSpec vs = {1, 0, 0};

    if (SLang_Num_Function_Args != 2 ||
        vec_pop( arg2_r, 0, 0, &vs) == -1 ||
        vec_pop( arg1_r, 0, 0, &vs) == -1 )
    { Slirp_usage(0,0,1); finalize_refs(VREF_2); return; }

    if (vec_validate(&vs, VREF_2) == -1) {finalize_refs(VREF_2); return;}
    VEC_ALLOC_RETVAL(double, VREF_2);
    for (_viter=0; _viter < vs.num_iters; _viter++) {
        retval[_viter] = hypot(*arg1,*arg2);
        VINCR_2;
    }
    VEC_RETURN(retval, 0, SLANG_DOUBLE_TYPE, SLang_push_double, 0, 1);
    finalize_refs(VREF_2);
}

```

(b)

```

static void sl_hypot (void)
{
    ...
    int _viter;
    VecSpec vs = {1, 2, 0};

    ...
    { Slirp_usage(0,0,3); finalize_refs(VREF_2); return; }
    ...

    #pragma omp parallel for
    for (_viter=0; _viter < vs.num_iters; _viter++) {
        retval[_viter] = hypot(arg1[_viter],arg2[_viter]);
    }
    ...
}

```

Fig. 7. Vectorized (a), and parallelized (b) wrappers for `hypot`.

they would be in a standard wrapper, with additional reference variables declared to record argument metadata such as array dimensions and stride. Two additional variables are declared to support vectorization: a scalar to index the vectorization loop, and a `VecSpec` structure to reflect the parameters of vectorization, which are adjusted by `vec_pop` as it marshals arguments from S-Lang. The `vec_validate` function ensures that `hypot` can safely be called with the given inputs; it may also adjust the stride of non-isomorphic arguments along the way and allocate space for arguments which have been omitted from the S-Lang wrapper invocation, such

```

define weibull_fit(lo, hi, params)
{
    variable a, b, x1, r=@lo, i, c, m, d, e;

    a = params[0]; b = params[1]; c = params[2]; d = params[3];

    x1 = c - b * ((a-1)/a)^(1/a);
    m = (hi+lo)/2.0;

    i = where(lo > x1);
    if(any(i))
        r[i] = d*(a/b)*((m[i]-x1)/b)^(a-1)*exp(-((m[i]-x1)/b)^a);

    i = where(lo <= x1);
    if(any(i))
        r[i] = 0;

    return r;
}

```

Fig. 8. The 4-parameter Weibull model in S-Lang, as a custom fit function for ISIS. The `lo` and `hi` arrays represent bin edges in a 1D grid; their sizes vary as the X axis value in Fig. 12.

as `double *result` from `vmult`, because they are still required by the wrapped function. The `VEC_ALLOC_RETVAL`, `VEC_RETURN`, `VREF_n`, and `VINCR_n` macros enhance readability by masking unnecessary detail, performing tasks like memory management, argument striding, and the expansion of argument lists for support routines.

Serial vectorization appeared first in SLIRP, but Fig. 7-(b) shows how easy it was to adapt for OpenMP: most of the parallel wrapper is identical to the serial version and has been elided. In addition to the OpenMP pragma – and the usage message which was changed to indicate that the wrapper is also parallelized, the loop index has been changed to a signed integer. This unfortunately reduces by half the maximum size of array operations which may be parallelized, and also led to more extensive code changes during the operator parallelizations described in §5.2, but was necessary for conformance with the OpenMP 2.5 specification; we look forward to the support for unsigned indices coming in OpenMP 3.0. Note that array indexing is used to locate elements, instead of pointer dereferencing. Pointers are used in serial wrappers for flexibility – they cleanly enable each argument to have its own stride. In OpenMP loops, however, pointer traversals introduce unwanted concerns for portability and possible non-conformance with the specification. Although techniques exist for iterating over pointers within OpenMP blocks [Massaioli et al., 2005], we avoid the additional complexity and potential performance degradation by using a single loop variable to index all arguments; the tradeoff is that all arguments must be isomorphic, enabling the same stride to be used for each.

5. PERFORMANCE EXPERIMENTS

In addition to `atof` and `hypot` wrappers were generated for the `sin`, `cos`, `exp` and `log` intrinsics. We chose to evaluate functions already available and vectorized in S-

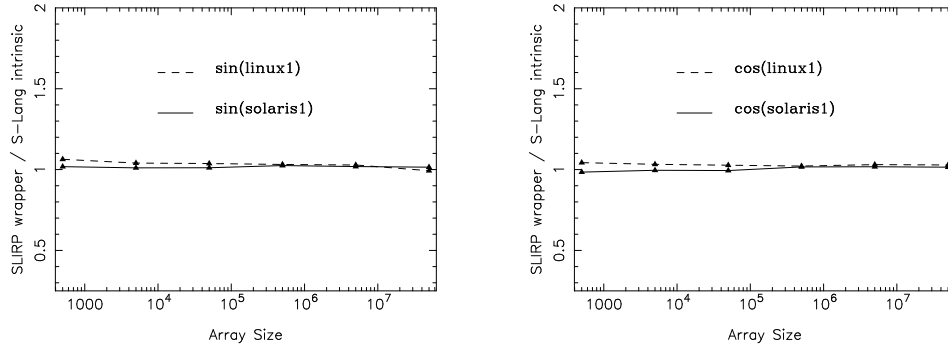


Fig. 9. Serial performance ratios of SLIRP-vectorized wrappers of `sin` and `cos` versus hand-crafted S-Lang intrinsics, on Linux1 and Solaris1. Mean runtimes per array size ranged from ca. $3.2\text{e-}5$ to 4.2 seconds on Linux1, and from $2.2\text{e-}4$ to 9.15 seconds on Solaris1.

Lang to illustrate several points: First, that SLIRP introduces minimal overhead; second, that S-Lang intrinsics can be replaced at runtime, without altering the semantics of numerical scripts⁷; and finally, that this leads to the use of existing serial algorithms in parallel contexts, without recoding. Consider for example the S-Lang function in Fig. 8, which defines a Weibull model [Weibull, 1939] for fitting in ISIS. While Fortran and C/C++ models may also be imported into ISIS, in just seconds with SLIRP, it can be faster to code them directly in S-Lang and avoid compilation steps during experimental tuning. The high performance numerics in S-Lang means that such interpreted models need not trade speed for convenience.

5.1 Parallel Functions

This model was taken from an active research project and originally coded for serial use; it was parallelized in two ways, neither of which involved changing a single line of its code. First, by using `import("par")` to dynamically load the module of parallel wrappers generated by SLIRP. This can be done either interactively at the command prompt (as in Fig. 4), programmatically in a script, or even automatically at ISIS launch by putting the `import()` within an `.isisrc` initialization file. The advantages of using SLIRP for parallelization are automation and simplicity: it may be employed immediately in S-Lang 2 applications by any user, merely by installing an OpenMP-aware compiler; having OpenMP support in GCC now makes this a much lower barrier than in the past.

5.2 Parallel Operators

A shortcoming of this tactic, however, is that it can only be used to parallelize functions, leaving a potential vulnerability to Amdahl's Law: in an expression such as `cos(x)/5 + sin(x)/2` the two divisions and one addition would still be computed serially. Our second approach to multiprocessing therefore involved manually parallelizing the `+`, `-`, `*`, `/`, `<=`, `>`, and `^` (exponentiation) operators by adding

```
#pragma omp parallel for if (size > omp_min_elements)
```

⁷With one exception: complex types are not supported by the wrappers parallelized here. SLIRP can wrap Fortran codes with complex arguments and return values, however.

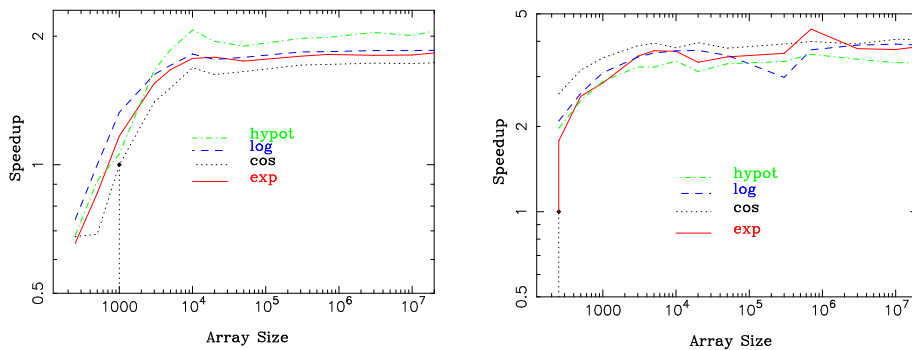


Fig. 10. Speedups from replacing selected S-Lang math intrinsics with parallelized versions generated by SLIRP -openmp. Left: Linux2. Right: Solaris4. Mean runtimes per array size ranged from ca. $1.7\text{e-}05$ to 2.6 seconds on Linux1, and from $3.8\text{e-}05$ to 9.1 seconds on Solaris1. The dotted vertical lines mark the inflection points where parallel performance begins to overtake serial, ca. 1000 elements on Linux2 and 250 elements on Solaris4.

to the operator loops as discussed in §1.2; the `where` function was partially parallelized, too, and utilized an additional `reduction` clause. Although parallel operators increase performance, a disadvantage of this approach is that it requires edits to the internals of S-Lang and these changes are not yet available to the general public. The `if` clause in the OpenMP directives was used to tune performance for small array sizes, where the cost of threads outweighs the serial execution time. During measurement the control variable was set with `getenv()` to one of the values `{0, 500, 1000, 5000, 10000, 50000, 100000}`.

5.3 Results and Analysis

Unless otherwise noted, the plots discussed here represent measurements of prerelease GCC 4.2 -O2 builds on Linux2 and Sun Studio 9 -xO3 builds on Solaris4, with position independent compilation. Comparable trends were seen in additional testing with the Intel 9.1 compiler on Linux2 and prerelease GCC 4.3 on an 2.33 Ghz Intel Core Duo Macintosh laptop (3 GB RAM) running OS/X 10.4.9. Runtimes were calculated by executing each function or operator 20 times per array size, and discarding the highest and lowest (with the exception of Weibull) before averaging.

The ratio plots in Fig. 9 indicate that the overhead of automatic vectorization in SLIRP is effectively constant and negligible: serial SLIRP wrappers are extremely competitive with the hand-crafted intrinsics in S-Lang, for both large and small arrays. Even without parallelism SLIRP vectorization therefore provides a solid path to increased performance when wrapping external codes. Purely serial vectorizations of `atof` and `strlen`, for instance, are ca. 50X faster than `array_map()`-ing the corresponding non-vectorized S-Lang intrinsics.

The speedup plots in Figs. 10, 11, & 12 demonstrate significant performance gains to be had from parallelism. Performance of the parallelized functions approaches the theoretical maximum of linear speedup as array sizes increase, and the inflection points in the size of the arrays needed for nominal speedup from multithreading (represented by the dotted vertical lines) are relatively small, ca. 1000 elements on Linux2 and 250 elements on Solaris4.

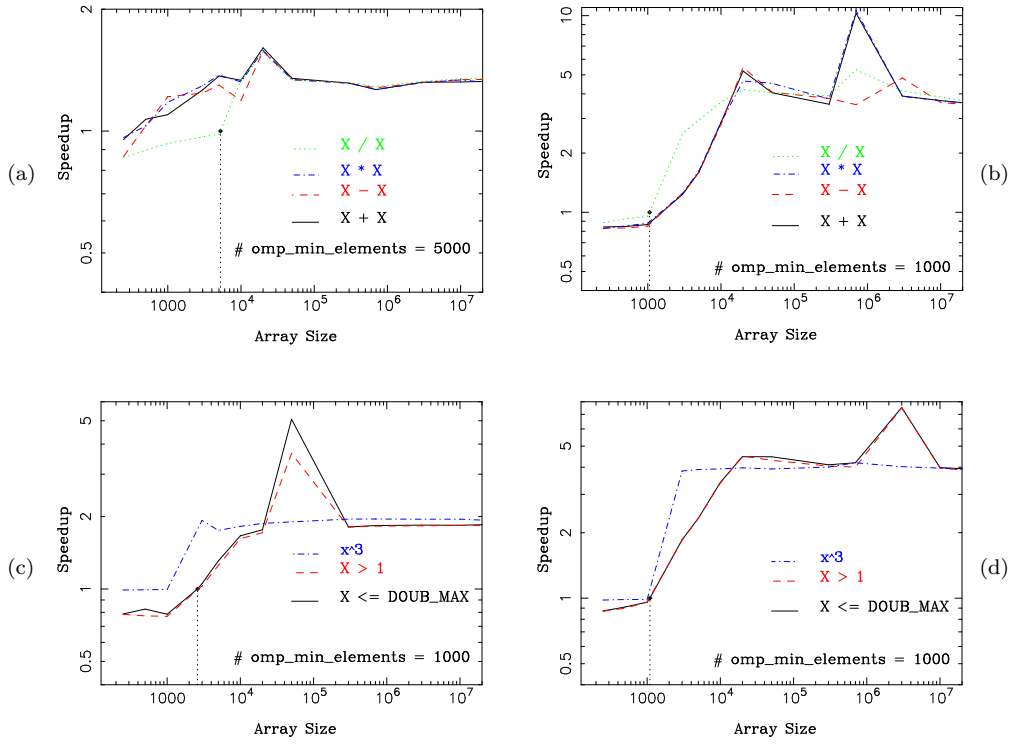


Fig. 11. Speedups from parallelizing selected S-Lang operators with OpenMP. Left Top: Linux2 with GCC 4.2. Left Bottom: Linux2 with Intel 9.1. Right: Solaris4. Linux2 inflection points are at 5246 and 2602 array elements, with mean runtimes from ca. $3.2\text{e-}6$ to 2.6 sec. Solaris4 inflection points are at 1059 and 1092 elements, with runtimes from ca. $1.1\text{e-}5$ to 13 sec. The superlinear spikes are discussed in §5.3.

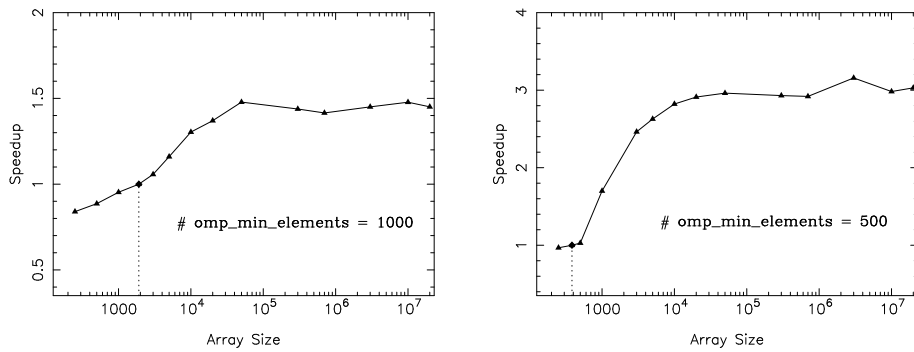


Fig. 12. Aggregate speedup of the Weibull fit function due to the parallelized operators and functions detailed above. Left: Linux2, with inflection point at 1907 array elements and mean runtimes from ca. $1.6\text{e-}4$ to 23 sec. Right: Solaris4, with inflection point at 384 elements and runtimes from ca $6\text{e-}4$ to 61 sec.

Fig.11-(a) shows that on Linux2 the core arithmetic operators did not parallelize as well as the functions, with speedups peaking at 60% before converging to 35%, although the gains are respectable. The lower speedup is not surprising: these operators can be executed directly as CPU instructions, requiring far less overhead than function invocation. On Linux2 these instructions execute fast enough to make the cost of parallel thread creation significant; on Solaris4 the same arithmetic instructions execute at a slower clock speed, resulting in greater speedups from parallelism. The relational and exponentiation operators converged to nearly linear speedups on both platforms. The excellent speedup of exponentiation stems from the operator being implemented in terms of the C `pow()` function, so we should expect its speedup curve to resemble those of Fig. 10. The relational operators parallelized well because they are not atomic CPU operations; they require many more assembly instructions to implement than, say, the division operator, approaching the number required for a short function call. The large superlinear spikes in Figs. 11 (b)-(d) appear consistently in every dataset collected. They do not reflect faulty parallelization of the S-Lang operators, because the serial and parallel results were verified identical and similar trends were observed with pure C codes written to perform the same computations. We attribute them to cache effects or page faulting that is less pronounced in parallel execution because each CPU receives a smaller portion of the problem.

In Fig. 12 these spikes are seen as smoother bumps at the corresponding array sizes. The Weibull model speedups converge on ca. 50% for Linux2 and 75% for Solaris4. While below the ideal of linear speedup, these are sizable performance increases; models with more calls to parallelized functions would exhibit even greater gains. These results have added significance in that end-users need to do nothing – in terms of learning parallelism or recoding sequential algorithms – to obtain them. Furthermore, recall that these models are used in the context of an iterative fitting process. Fits do not converge after just one iteration, and generating accurate confidence intervals – an absolute necessity for credible modeling – can require that thousands of fits be performed at each point on a parameter space grid, with potentially a million or more fits performed for a single pair of parameters, and tens of millions if multiple parameter sets are to be explored. In such cases the speedups given here accumulate to significant differences in the overall runtime of an analysis sequence. By transparently using OpenMP to effect greater multiprocessor utilization we gain the freedom to explore on the desktop more challenging problems that other researchers might avoid for their prohibitive cost of computation.

6. CONCLUSION

Multicore chip designs are making it possible for general users to access many processors. At the granularity of the operating system it will be relatively easy to make use of these extra cores, say by assigning whole programs to separate CPUs. As noted with increasing frequency of late, though, it is not as straightforward to exploit this concurrency within individual desktop applications. In this paper we demonstrated how we have helped our research colleagues prepare for this eventuality. We have enhanced the vectorization capabilities of SLIRP, a module generator for the S-Lang numerical scripting language, so that wrappers may be annotated

for automatic parallelization with OpenMP. This lets S-Lang intrinsic functions be replaced with parallelized versions, at runtime, without modifying a single line of internal S-Lang source. We have shown how S-Lang operators may also be parallelized with relative ease, by identifying key loops within the interpreter source, tagging them with OpenMP directives and recompiling. These simple adaptations have yielded beneficial speedups for computations actively used in astrophysical research, and allow the same numerical scripts to be used for both serial and parallel execution – minimizing two traditional barriers to the use of parallelism by non-specialists: learning how to program for concurrency and recasting sequential algorithms in parallel form. We extrapolate that the advent of widespread OpenMP support in free compilers such as GCC presages a proliferation of multicore-enabled scientific codes in the open source community, parallelized in largely the manner given here.

ACKNOWLEDGMENTS

This work was supported by NASA through the AISRP grant NNG06GE58G (HYDRA) and Smithsonian Astrophysical Observatory contract SV3-73016 for the Chandra X-Ray Center. The author would like to thank his MIT colleagues for thoughtful review and constructive criticism.

REFERENCES

- Beazley, D. (1996). SWIG : An Easy to Use Tool for Integrating Scripting Languages with C and C++. In *4th Tcl/Tk Workshop Proceedings*.
- Chuang, W., Narayanasamy, S., Venkatesh, G., Sampson, J., Biesbrouck, M. V., Pokam, G., Calder, B., and Colavin, O. (2006). Unbounded page-based transactional memory. In *ASPLOS-XII: Proceedings of the 12th international conference on Architectural support for programming languages and operating systems*, pages 347–358, New York, NY, USA. ACM Press.
- Creeger, M. (2005). Multicore CPUs For The Masses. *Queue*, 3(7):64–ff.
- Creel, M. (2005). User-Friendly Parallel Computations with Econometric Examples. *Comput. Econ.*, 26(2):107–128.
- Folk, M., McGrath, R., and Yeager, N. (1999). HDF: an update and future directions. In *IEEE Geoscience and Remote Sensing Symposium, 1999, Vol. 1*, pages 273–275.
- Houck, J. C. (2002). ISIS: The Interactive Spectral Interpretation System. In Branduardi-Raymont, G., editor, *High Resolution X-ray Spectroscopy with XMM-Newton and Chandra*.
- Karlsson, S. and Brorsson, M. (2004). A Free OpenMP Compiler and Run-Time Library Infrastructure for Research on Shared Memory Parallel Computing. In *Proceedings of The 16th IASTED International Conference on Parallel and Distributed Computing and Systems*, Calgary, AB, Canada. Acta Press.
- Kusano, K., Satoh, S., and Sato, M. (2000). Performance Evaluation of the Omni OpenMP Compiler. In *ISHPC '00: Proceedings of the Third International Symposium on High Performance Computing*, pages 403–414, London, UK. Springer-Verlag.
- Lee, E. A. (2006). The Problem with Threads. *Computer*, 39(5):33–42.
- Massaioli, F., Castiglione, F., and Bernaschi, M. (2005). OpenMP Parallelization of Agent-Based Models. *Parallel Comput.*, 31(10-12):1066–1081.
- Noble, M. S., Houck, J. C., Davis, J. E., Young, A., and Nowak, M. (2006). Using the Parallel Virtual Machine for Everyday Analysis. In Gabriel, C., Arviset, C., Ponz, D., and Enrique, S., editors, *ASP Conf. Ser. 351: Astronomical Data Analysis Software and Systems XV*, pages 481–+.
- Weibull, W. (1939). A Statistical Theory of the Strength of Materials. *Ingenior Ventenskaps Akademien Handlingar*, 151.